



IBM Enhances Storwize V3700 Again — Increasing Storage Capabilities for the Smaller Data Center

Analyst: David Reine

Management Summary

Every day thousands of shoppers go to a car dealership looking for just the *right car* at just the *right price*. These dealers can show you everything from the smallest, economy vehicle to the most expensive luxury model, with a range of mileage performance and a multitude of features. Most likely, you do not need to look elsewhere, as every dealer has everything that you could possibly want – although you may have to settle for *something less*, less mileage or fewer features, as your automotive budget may not be sufficient to acquire all that you desire. You may have to compromise to fit your budget. You may not be able to afford the top of the line sound system, or the navigation, or some other feature that is offered and that you really wanted. As an important rule of thumb, you cannot get the luxury model at an economy price.

The idea of having to settle for something less is nothing new to the managers of data centers around the globe. Take storage, for example. The data center has many specific needs to be satisfied: capacity, scalability, ease-of-use, lowest possible total cost of ownership (TCO), and performance to name a few. Unfortunately, the IT budget often will not support the acquisition of the platform that provides virtually unlimited capacity, the highest performance, and has the least complexity. Fortunately, *your* smaller data center may not need to support exabytes of storage or hundreds of thousands of IOPS. You only may need to find the platform with the highest scalability and highest performance that meets your business requirements and fits within the IT budget. *What features might you crave that until recently have only been available on high-end models?* Your list might include a multi-tier environment including SSD and HDD storage, automated data migration between tiers, replication, thin provisioning, ease of use, and more automation to decrease operational complexity, and many more features to improve administrative efficiency.

The largest enterprises have the budget to afford exabyte capacity and full functionality, maybe including even heated seats and a GPS. That is not always the case for the rapidly growing, smaller business or government agency. Fortunately, that no longer is a problem with IBM's *Storwize V3700*. First announced in 2012, the *Storwize V3700* represents a smaller version of IBM's *Storwize V7000*. Now, IBM has made it even better, enhancing the *V3700's* capacity and upgrading it with most of the same software as on other, more expensive *Storwize* systems. With the *V3700* the smaller data center does not have to settle for *something less*; it can acquire *something more*, with the capacity needed for today and next year *and* with the functionality that they need, all while staying within budget. To learn more about IBM's upgraded *Storwize V3700*, please read on.

Storage Growth Pains

Today's smarter data center is built with systems that increasingly are scalable, instrumented, interconnected, and intelligent. They have to meet the challenges facing businesses and agencies that need to stay responsive in a perpetually changing

IN THIS ISSUE

➤ Storage Growth Pains	1
➤ IBM Storwize V3700	4
➤ IBM Storwize Family Software v7.2	6
➤ Conclusion	7

environment. Usually, the business has a growing volume of data with increasing value, but storing it can be expensive and very difficult to manage, both due to the complexity of the architecture and the increasingly-rapid growth of information. Just about all IT organizations have budgetary constraints and cannot afford to make investment mistakes, even more so for the smaller data center with a smaller budget. While this smaller data center may require an entry-level storage solution, it still requires some enterprise-class functionality. It has to manage its investment in storage wisely in order to maintain budgetary control. For a smaller data center, a mistake can have damaging effects, potentially jeopardizing operational viability and possibly endangering business continuity and sustainability. There is one more reality for a 2014 data center. Storage is growing even faster than in 2013 and, with little or no restraint throughout your business or agency, it might eat through your entire IT infrastructure budget. If unchecked, storage tends to try to consume every dollar it can find. It does not matter what size data center you have, your storage capacity requirements probably are growing at a rate that seems to be without limits.

Needing Multiple Tiers of Storage

Data center managers always are looking for higher performing storage with more reliability for their critical applications, while seeking high-capacity, low-cost storage for less critical applications. In addition, they seek more capacity for replication and data protection processes. Over-provisioning of LUNs for what might possibly be needed in the future usually results in significant waste today, due to wasted space reserved for future use. This is a common practice, as no storage administrator wants to have short headroom, which would require constant attention and modification. So there is a strong tendency to over-allocate. Typically, it is only near the end of its useful life that a storage solution is at its fullest utilization. This can have major implications for your data center budget.

What is needed is a multi-tiered storage environment spanning solid state storage (SSDs) at Tier 0, for the fastest response possible; high-performance drives for the typical mission-critical applications using Fibre Channel (FC) or SAS connectivity at Tier 1; and high-capacity storage, usually satisfied with high capacity SATA or SAS drives at Tier 2. In case you lost count, that is three tiers of storage, which might optimize your storage budget but also might strain or stretch

beyond the management abilities and time availability of your storage administrators.

In the largest data centers, storage area networks (SANs) connected to islands of special-purpose arrays often also are competing for increasingly expensive administrative attention with islands of high-capacity lower-cost arrays, as administrators try to determine which array is the correct resource for storing an application's needed data. This chore is compounded massively by the realities of real-world data and business activities, which tend to change in both expected and unexpected ways. A business promotion (like a "two-day sale") almost certainly will cause a spike in business activity, online access, and transaction processing. This is to be expected.

However, often, the data center doesn't know in advance what is going to be "hot", whether it is data or processing. This is a "chicken-or-egg" kind of situation. *Is the data hot because an application is hot, or is it because there is something hot about the data being sought?* Regardless, the data is hot and performance will lag if it is not served with alacrity. Trying to adjust assets and storage tier allocations in response to this in real time is very difficult, time consuming, and often foolhardy, even for larger data centers. Obviously, it is nearly-impossible for smaller data centers to manage this successfully. Additionally, different applications may need access to a range of storage resources, e.g., database-driven applications may be competing for storage resources also being used to access unstructured objects and files. If they all are treated the same, the application mission-critical (time-sensitive) application may not perform satisfactorily. This all leads to a most-important question: ***How many storage solutions do you need to satisfy the many applications used by your business?*** Having a wide array of storage solutions is a luxury few small data centers can afford.

Minimizing Storage Networking Pain

Additionally, today's smaller data centers need to provide the same seamless delivery of IT services as do larger enterprise data centers, perhaps without the complexities of FC. Hopefully, this can be accomplished via a single storage architecture that simplifies the support of all critical applications and data types, regardless of their server environment – physical, virtual, or cloud.

Meeting Important Data Requirements

Today, many modern businesses probably are mining their data for more valuable information

regularly, and often needing to do so in short order. With an increasing number of processes being done in real time, there are a growing number of challenges to every data center that needs to stay responsive in a dynamic environment. Because applications, such as data mining, now are becoming critical to the operational success of the business, the managers of many data centers are demanding that the deployed storage provide continuous availability to all of the data while also preserving the IT budget.

Applications that once seemed to be less than mission-critical (end-of-the month analytics, for example) have become mission-critical because the time frame has become near real-time, instead of being run at night or weekends after the close of the reporting period. Other applications, such as email, which in the past may have been deemed to be business-critical, may now be mission-critical because it is part of the regular flow of business.

On the server front lines, application consolidation and server virtualization have enabled the IT staff of smaller data centers to reduce the TCO attributable to the server infrastructure. This was accomplished by centralizing server resources and supporting multiple applications on a single server, thereby limiting the amount of staff required to manage these systems and also limiting the number of software licenses required. This leads to another important question: ***What is being done to reduce the TCO for storage in the smaller data center?***

A virtualized server environment, with networking being shared among many applications, usually places many different demands on storage than single-purpose servers, where each is connected separately. This can mean that the services being delivered may be impaired operationally by applications in other virtualized partitions on the same physical server. This tends to create new challenges, including requiring higher I/O throughput, and new problems, such as new hot spots and the resulting need for migration to a more (or less) responsive tier. Collectively, these likely will require either putting all of the storage on a tier-1 storage solution (an expensive “set and forget” strategy), which is very expensive, or by having sufficient automated intelligence to handle multiple tiers of storage dynamically. With multiple applications sharing the same server resources, the smaller data center now must look toward upgrading its storage infrastructure to support storage virtualization and multiple tiers of storage,

from high-capacity HDDs¹ to high-speed SSDs², to provide scalability, and to accommodate a virtualized environment.

Connectivity Needs

Improving storage performance in larger enterprise data centers usually means upgrading high-performance storage to support increased throughput, while ensuring continuous access through high availability. For the smaller data center, this usually means not being forced unnecessarily to change the way servers are connected to storage and simplifying the network connectivity with sufficient paths and enough ports to avoid the high acquisition costs and higher administrative costs usually associated with FC³ SANs. Any viable, smaller data center storage solution needs to be able to deliver, both interconnection convenience and performance. To accomplish this, many smaller data centers are turning to iSCSI over Ethernet for their storage networking protocol and some also are upgrading to 10Gb Ethernet to simplify their networking infrastructure.

Simplification Needs

Smaller data centers also may want to simplify administration by consolidating all existing drive pools onto a single array, possibly simplifying connectivity by moving all storage access to a sharable network, and also by getting the benefit of common management via a suite of integrated management software. This will enable the data center staff to use a less expensive mix of tiered storage, and to use a simpler and more feature-rich administrative toolset.

Using yesterday’s storage solutions and multiple storage architectures to solve today’s data center challenges certainly is not efficient, especially for the smaller data center, which now probably needs to change its outmoded storage paradigm, possibly by eliminating multiple silos of storage currently populating the data center and consolidating onto a single, virtualized array. By evolving from multiple silos of storage to a single, consolidated storage platform, the data center can simplify the environments and lower administrative costs with easy-to-use tools. This unification also can prepare the way for the delivery of ITaaS⁴ and cloud-based storage.

¹ HDDs=hard disk drives.

² SSDs=solid-state disks.

³ FC=Fibre Channel.

⁴ Information Technology as a Service.

Today, the needs and desires of smaller data centers much more closely match with those of larger data centers. In the past, the smaller data center might have been more driven by what it could afford. Today, it is based on what it needs to succeed, including simplicity, ease of management, lower TCO per GB, etc. Usually, the smaller data center needs most of the storage functionality that is built into the hardware and software features found on enterprise-class storage arrays.

Management Software Needs

The technology requirements for storage management software for the smaller data center include much of the same functionality provided to many high-end arrays in larger data centers. These include:

- Easy-to use storage management;
- Automated data retention and protection;
- Thin provisioning for improved storage utilization⁵; and
- Storage services, including performance monitoring and reporting.

Hardware Needs

In terms of hardware, that data center is looking for a single multi-tiered, storage solution that includes most, if not all, of the following features found in larger enterprise arrays.

- **Solid State Drives (SSDs)** at Tier 0, with very high IOPS, for the critical applications requiring highest performance. This often is the case for accessing and updating databases.
- **High-performance SAS drives**⁶ at Tier 1 to satisfy mission-critical application performance requirements at a lower cost than with SSDs or FC drives.
- **High-capacity drives**⁷ at Tier 2 to fulfill the rapidly growing needs for less-critical data, such as backup, email, and, website data. This enables the data center to put less critical data on lower-cost drives.
- **Data integrity features** to ensure the reliability of the data. These include dual-ported disk

drives with automatic failure detection and RAID rebuild capability with global hot spares.

- **Intelligent scalability** to improve data utilization to enable RAID arrays that span multiple chassis.
- **Internal virtualization** capabilities, like thin provisioning, to improve the capacity utilization rates of the storage array, in much the same way that the data center has consolidated and virtualized its server infrastructure to improve its server TCO. This also makes management simpler, because you don't have to manage the physical mappings for each LUN.

Finding the Right Storage Solution

Many storage vendors have a product set that they can try to force-fit into any solution. IBM is different; it has an extremely broad range of storage solutions for the enterprise data center, including the *DS8000*, *XIV*, and *Storwize V7000*⁸. However, rather than trying to force-fit one of these into a smaller data center, IBM chose to meet the specific needs of the smaller data center by scaling down the IBM Storwize V7000 with the announcement in November 2012 of the IBM Storwize V3700, a little brother of the high-end V7000.⁹ This isn't what you might be thinking. The V3700 isn't a stripped down, entry-level product that is only similar in naming to the V7000. As you will see shortly, it has the features and software that the small data center needs, all of which are the same as on the V7000.

Yes, some of the V7000's advanced features, such as compression, have not been enabled to keep the costs and complexity down. However, having the same tool set to administer both small and large arrays makes life much simpler for the data center staff, especially if managing more than one array or when it comes time to move up to a V5000 or V7000. This singularity means that the V3700 has the right stuff. Let's look at the details.

The IBM Storwize V3700

The V3700 is an entry-level array only in the sense that it is modestly priced at \$11K for low-


⁵ Thin Provisioning optimizes efficiency by allocating disk storage space in a flexible manner among multiple users, based upon the minimum space required by each user at any given time. In effect, the applications think that they have plenty of headroom but, in reality, it is added transparently as it is needed.

⁶ High-performance drives rotate at 10,000 or 15,000 RPM.

⁷ High-capacity drives rotate at 7200 RPM or slower.

⁸ See [The Clipper Group Navigator](#) entitled *IBM Brings Enterprise Functionality to Mid-Range Storage* dated October 7, 2010, and available at <http://www.clipper.com/research/TCG2010047.pdf>.

⁹ For more on the initial announcement of the IBM V3700, see [The Clipper Group Navigator](#) entitled *IBM Let's You Have It Your Way – Introduces an Enhanced Storwize V3700* dated June 10, 2013, and available at <http://www.clipper.com/research/TCG2013010.pdf>.



end capacities and capabilities. (See Exhibit 1, above.) As you will see, it has a broad range of scale and thus should be able to suit the needs of most smaller data centers looking for block storage. It will enable the elimination of existing islands of storage and facilitate the migration of external storage in to the V3700. With the V3700, the data center staff gets most of the same large enterprise data center functionality as in the V7000, in a platform with a little less scalability (about 50% less maximum capacity) and less total performance. With a maximum capacity of 240 terabytes, this represents more capacity than most smaller data centers need today, or will need tomorrow. The V3700 enables the smaller data center with the capability to set its expectations to its requirements and ambitions, scaling down capacity without seriously scaling down functionality and needed performance.

Originally announced in 2012, today's Storwize V3700 has seen multiple upgrades to both capacity and functionality. The V3700 consists of a 2U, 19-inch rack-mount control enclosure and up to four 2U expansion drawers. The control unit is configured with two dual, active-active control canisters for improved reliability and load sharing. Each canister comes with 4GBs of cache, for a total of 8GBs. IBM also provides an optional upgrade to 16GBs of cache, in total. Each control unit is configured with four 1Gb iSCSI host ports and six 6Gb SAS host ports, for consolidation

flexibility. In addition each control enclosure can be configured with eight optional 1Gb iSCSI ports, eight optional 6Gb SAS ports, eight optional 8Gb FC ports, or four optional 10Gb iSCSI/FCoE ports for even more flexibility. As a result, the V3700 can support more host I/O than the typical smaller data center might consume and can handle your choice of protocols.

IBM offers two control enclosures, one supporting (12) 3.5" LFF¹⁰ drives, the other supporting (24) 2.5" SFF¹¹ drives. The SFF enclosure supports a mix of both HDDs and SSDs. The V3700 control enclosure supports up to four 2U expansion enclosures, each with dual-active expansion canisters, with each enclosure capable of supporting up to 24 SFF drives or 12 LFF drives, for a total of 120 SFF or 60 LFF drives. The expansion drawer containing the 2.5" SFF drives can mix both HDDs and SSDs in the enclosure. The data center staff can also mix drive capacities in an enclosure, although they all must be of the same form factor.

Since the 2012 announcement, the product set

¹⁰ LFF=Large Form Factor. LFF drives also are called 3.5" drives. Currently, the maximum capacity of a LFF drive is 4 TBs.

¹¹ SFF=Small Form Factor. SFF drives also are called 2.5" drives. Currently, the maximum capacity of a SFF drive is 1.2 TB. Solid-state drives (SSDs) used in arrays tend to be SFF drives.

has been expanded to support up to 1.2TB SFF HDDs and up to 800GB SSDs¹². In terms of LFF HDDs, the data center has an even wider range, with up to 4TB at 7200 RPM, up to 1.2TB at 10K RPM, and 300GB devices at 15K RPM. In terms of storage capacity, the V3700 has the capability, using all four expansion chassis, to scale to a maximum capacity of 240TB when using 4TB LFF drives, with up to 16 GB of cache in an active-active configuration with dual controllers.

Each V3700 enclosure also has redundant, hot-swappable power supplies and fans and supports RAID levels 0, 1, 5, 6, and 10, for improved reliability. The LFF enclosure has been designed to provide high capacity with low-cost nearline drives, while the SFF enclosure is designed for high-performance drives and SSDs, as well as high-capacity, nearline drives.

The Storwize V3700 is certified to meet the Telco industry's specification for both AC and DC models, the latter to be used where AC power is not available. The DC model uses a -48V DC power source and only is available with SFF drives. The Telco models are compliant with both NEBS (Network Equipment-Building System) Level 3 and ETSI (European Telecommunications Standards Institute) standards for reliable operation.

IBM Storwize Family Software v7.2

These days, there's a lot more to storage than hardware. A rich software set is what likely will separate the leader from the less desirable. So let's look at the functionality that IBM has included with the Storwize V3700 that also is provided in the high-end Storwize models.

- **Data Migration** enables a non-disruptive migration from existing legacy SAS storage to Storwize V3700 and comes standard with the V3700.
- **Virtualization** of internal storage enables rapid, flexible application deployment, to simplify and standardize application management, and enable configuration changes to meet changing business needs. It also enables the data center to transform the economics of storage with consistent data efficiency, virtually eliminating application downtime for data movement.
- **Thin Provisioning** enables applications to consume only the space they are actually

using, not the total space that has been allocated to them.


- **Turbo Performance**¹³ increases the maximum IOPS of the V3700 by 50% and throughput by 100%. It is a licensed function that provides enhanced performance for the system. *When might this be worthwhile?* Configurations with greater than 80 disk drives or more than five SSDs are likely to benefit from the increased IOPS offered with Turbo performance. Configurations with greater than 30 disk drives are likely to benefit from the throughput increase offered with Turbo performance. This option does not require hardware changes and is a non-disruptive upgrade.
- **IP replication** via the integrated *Bridgeworks SANslide* network optimization, provides a lower cost option for deploying remote replication, using artificial intelligence to improve network bandwidth utilization by a factor of three times. It enables the use of Ethernet connections for optional remote mirroring with no separate appliance required, and is transparent to both servers and applications. Faster replication cycles result in better remote data currency and faster recovery. Local replication is included standard in v7.2.
- **FlashCopy** allows the user to create copies of data for backup, parallel processing, testing, and development. The V3700 supports 64 targets per system, or up to 2040 with the FlashCopy upgrade option
- **IBM System Storage Easy Tier** automates data placement onto two tiers of storage and to intelligently align the system with current workload requirements for one tier on SSD and the other on HDD. Initially developed for the *DS8700*, V3700 Easy Tier leverages a tier of solid-state storage in addition to a tier of HDDs.¹⁴ Easy Tier manages the solid-state storage as a large and low latency cache for the "hottest" data, while preserving advanced disk system functions, such as RAID protection and remote mirroring. **You don't have to manage the migration between the tiers, as Easy Tier does this for you, as the data**

¹² IBM uses enterprise-class, multi-level cell (E-MLC) SSDs.

¹³ The Turbo Performance Option only is available on the V3700.

¹⁴ While only two tiers are supported by Easy Tier and one must be SSDs, you can have different pairings of SSDs and HDDs in each SFF enclosure, as long as they are in separate enclosures.

Exhibit 2 — IBM Easy Tier – Automated Storage Tiering



- Constantly analyzes I/O and migrates data to **optimize performance**
- **Set it and forget it** ... unlike competitive tiering offerings

Improve performance
up to **3x**

... using as little as **5%**
flash storage*

* Audited Storage Performance Council benchmark

Source: IBM

heats and cools. (See Exhibit 2, above.)

- **An application-aware storage API** helps to deploy storage more efficiently by enabling applications and middleware to direct more optimal placement of data by communicating important information about current workload activity and application performance requirements.
- **A simple and easy-to-use GUI** to enable storage to be quickly deployed and easily managed, the same as for the V7000.
- **IBM Storage Mobile Dashboard** to provide monitoring and health check for Storwize systems from mobile devices.
- **VMware 5.5 and VASA block support**, in order to retain currency with the latest VMware capabilities.

Conclusion

While all data centers may not have been created equal, every data center staffer expects to have a rich set of features and functionality to do *their* job. With a “limited” capacity, if you want to call 60 LFF drives with 240 TB of capacity limiting, the V3700 provides the smaller data center with all of the block data functionality that it requires, from data migration to multi-tiering, and from virtualization to thin provisioning.

IBM has configured the new Storwize

V3700 specifically to handle the block storage needs of the smaller data center. Whether you need to connect via iSCSI, SAS, FC, or FCoE, the V3700 provides the smaller data center with the all of the storage it needs, no make that all of the highly reliable and efficient storage that it needs. **Derived from the same software set that runs the Storwize family systems, the V3700 delivers best-in-class features at an economy price.** SSDs, a large cache, and Easy Tier all contribute to provide the smaller data center with larger enterprise performance and administrative experience.

The Storwize V3700 enables the smaller enterprise to start small and grow at its own pace, investing what is needed, when it is needed, all while retaining control of the IT infrastructure budget. If you find yourself caught between the scalability “rock” and the budgetary “hard place”, you need to look into IBM’s Storwize V3700 for both today’s and tomorrow’s storage requirements.



About The Clipper Group, Inc.

The Clipper Group, Inc., now in its twenty-second year, is an independent publishing and consulting firm specializing in acquisition decisions and strategic advice regarding complex, enterprise-class information technologies. Our team of industry professionals averages more than 40 years of real-world experience. A team of staff consultants augments our capabilities, with significant experience across a broad spectrum of applications and environments.

- ***The Clipper Group can be reached at 781-235-0085 and found on the web at www.clipper.com.***

About the Author

David Reine is a Senior Contributing Analyst for The Clipper Group. Mr. Reine specializes in enterprise servers, storage, and software, strategic business solutions, and trends in open systems architectures. In 2002, he joined The Clipper Group after three decades in server and storage product marketing and program management for Groupe Bull, Zenith Data Systems, and Honeywell Information Systems. Mr. Reine earned a Bachelor of Arts degree from Tufts University, and an MBA from Northeastern University.

- ***Reach David Reine via e-mail at dave.reine@clipper.com or at 781-235-0085 Ext. 123. (Please dial “123” when you hear the automated attendant.)***

Regarding Trademarks and Service Marks

The Clipper Group Navigator, The Clipper Group Explorer, The Clipper Group Observer, The Clipper Group Captain's Log, The Clipper Group Voyager, Clipper Notes, The Clipper Group Calculator, and “clipper.com” are trademarks of The Clipper Group, Inc., and the clipper ship drawings, “*Navigating Information Technology Horizons*”, and “*teraproductivity*” are service marks of The Clipper Group, Inc. The Clipper Group, Inc., reserves all rights regarding its trademarks and service marks. All other trademarks, etc., belong to their respective owners.

Disclosures

Officers and/or employees of The Clipper Group may own as individuals, directly or indirectly, shares in one or more companies discussed in this bulletin. Company policy prohibits any officer or employee from holding more than one percent of the outstanding shares of any company covered by The Clipper Group. The Clipper Group, Inc., has no such equity holdings.

After publication of a bulletin on *clipper.com*, The Clipper Group offers all vendors and users the opportunity to license its publications for a fee, since linking to Clipper’s web pages, posting of Clipper documents on other’s websites, and printing of hard-copy reprints is not allowed without payment of related fee(s). Less than half of our publications are licensed in this way. In addition, analysts regularly receive briefings from many vendors. Occasionally, Clipper analysts’ travel and/or lodging expenses and/or conference fees have been subsidized by a vendor, in order to participate in briefings. The Clipper Group does not charge any professional fees to participate in these information-gathering events. In addition, some vendors sometime provide binders, USB drives containing presentations, and other conference-related paraphernalia to Clipper’s analysts.

Regarding the Information in this Issue

The Clipper Group believes the information included in this report to be accurate. Data has been received from a variety of sources, which we believe to be reliable, including manufacturers, distributors, or users of the products discussed herein. The Clipper Group, Inc., cannot be held responsible for any consequential damages resulting from the application of information or opinions contained in this report.